Using Agent-Based Technology to Create a Cost Effective, Integrated, Multimedia View of the Electronic Medical Record

Henry J. Lowe M.D., William K. Walker BS, John K. Vries M.D.
University of Pittsburgh, Pittsburgh PA

ABSTRACT

Image Engine is multi-user, client-server database for the storage, retrieval and sharing of a wide range of digitized biomedical images under development at the University of Pittsburgh. This paper provides an overview of the system and describes the use of agent-based technology to integrate clinical information from the Image Engine database and the MARS clinical information system at the University of Pittsburgh Medical Center. Agent-mediated links provide a mechanism for combining clinical data from multiple databases to create a unified, multimedia view of the electronic medical record.

INTRODUCTION

The electronic medical record is inherently a multimedia entity, incorporating clinical data in the form of text, images and sound. Most electronic medical record systems store only the textual component of this data resulting in a fragmented clinical database. While the textual reports for diagnostic images may be available electronically the images themselves are often difficult for the physician to access and integrate with other relevant clinical data[1]. The solution to this problem depends, in part, on the development of innovative medical multimedia database models that can integrate clinical data from multiple sources. Such systems may improve the quality of patient care[2], increase the patient's involvement in clinical decision making[3] and may produce significant new medical knowledge[4].

Systems for the storage, retrieval and classification of digital images are in their infancy[5]. Traditional Picture Archival and Communications Systems (PACS)[6] are generally expensive[7] , monolithic systems which serve primarily the needs of radiologists and are often not well integrated with other clinical databases. Integrated multimedia clinical information systems such as the VA's DHCP system[8] are rare and expensive to implement.

The relative immaturity of multimedia database technology along with the significant capital investment in currently installed text-based clinical information systems requires a cautious, stepwise migration towards large scale integrated multimedia clinical databases. Recent innovations in hardware, software and networking make it technologically and financially feasible to begin prototyping multimedia clinical information systems that function in parallel with an existing text-based electronic medical record. This paper describes such a system, called Image Engine™, under development at the University of Pittsburgh Medical Center (UPMC) as part of the National Library of Medicine's (NLM) Biomedical Applications of High Performance Computing and Communications (HPCC) initiative.

IMAGE ENGINE

Image Engine[9] is an object-oriented, multi-user, client-server database system for the storage, retrieval and sharing of a wide range of digitized clinical images. Image Engine uses agent-based technology to automatically and dynamically link clinical images with textual data in UPMC's MARS[10] clinical information system. The philosophy behind this project is to prototype a low-cost, expandable clinical multimedia database system that can seamlessly integrate both image and textual information from a variety of sources and provide a unified, multimedia view of the patient record. Image Engine is intended to function in parallel with existing clinical information systems to achieve this goal.

Image Engine consists of four inter-related components: (1) a central Object Database server; (2) multiple image file servers; (3) a graphical client application and (4) software agents that retrieve data from external database systems.

1. The Object Database

Each image stored in the Image Engine system is represented in a central Object Database as a set of linked entities representing patient, procedure and image properties. The current prototype uses a hybrid relational/hierarchical database architecture that supports a true object-oriented data model at the client level. This data model views images as being complex objects generated by clinical procedures performed upon patients.

0195-4210/95/$5.00

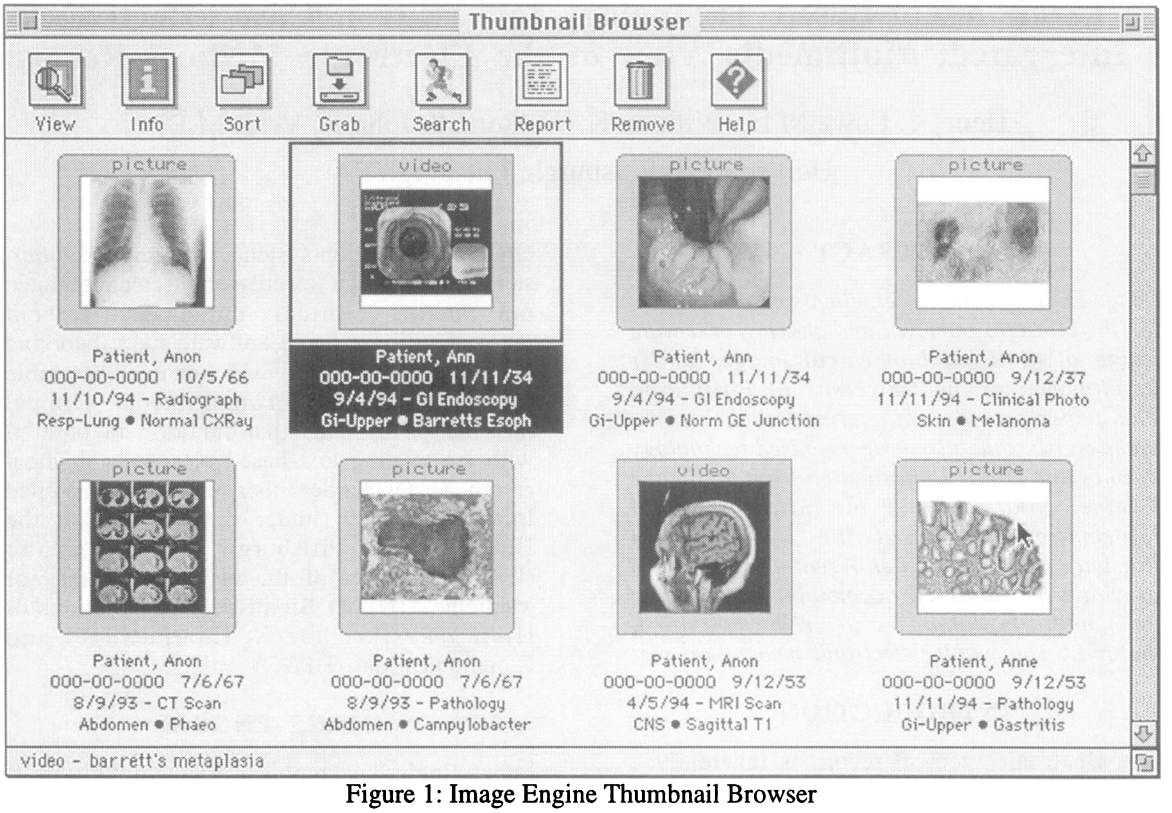

Figure 1: Image Engine Thumbnail Browser

In addition to storing properties such as patient name, procedure type, anatomic location etc. each image is also represented within the Object Database as a 100 x 100 pixel scaled "thumbnail" representation of the actual image stored on an image server. Almost all client-server database queries are handled by the Object Database.

2. Image Servers

Each image represented in the Object Database is stored on a network-based image server. These servers store still images as JPEG[11] compressed files and digital video images as either MPEG[12] or QuickTime™ compressed files. Each image object in the Object Database knows where its image file is stored and cooperates with the Image Engine client application to retrieve that image file if requested. The current image server prototype is based on an Apple Power Macintosh 8100/100 with 24 gigabytes of RAID (Redundant Array of Inexpensive Disks) level four disk storage. Multiple network-based image servers are supported potentially providing a very large image storage space using relatively inexpensive hardware and software. Our current image server configuration provides image storage at a cost of approximately 90¢ per megabyte. Image compression effectively increases the storage capacity of these servers by a factor of 10-20.

3. The Image Engine Client

The user interacts with both the Object Database and image servers via a graphical client application. The current Image Engine client prototype runs on Apple Power Macintosh computers with 40 megabytes of RAM, accelerated display hardware and 20 inch 1152 x 870 pixel resolution color displays. The client application allows one to create and execute complex database queries using a graphical Query Editor. Successful searches result in the creation of a Thumbnail Browser [Figure 1] displaying 100 x 100 pixel representations of images along with data identifying the patient, procedure and image type. Images may be retrieved from image servers, displayed, processed, scaled, sorted, copied and annotated using this browser. The client application can display 8 bit color or gray scale as well as 16 and 24 bit high resolution color images. Multiple images can be viewed simultaneously in separate, scrollable windows. Digital video images can be viewed and controlled using an on-screen VCR-like interface. Still frames can be extracted from digital video images. The Image Engine client can dynamically link with external applications to perform tasks such as complex image processing, sending images and data via electronic mail and thesaurus browsing.

4. Image Engine Agents

Image Engine was designed from the outset to function in parallel with existing, external clinical databases. While the system can serve as a stand-alone medical image database system, our model is to have Image Engine store only clinical images and the data required to permit efficient retrieval of these images. Additional related data such as procedure reports, laboratory data etc. are assumed to be stored in an external, text-based clinical information system. Image Engine will retrieve and display this information from external databases in real-time on an "as needed" basis. This model would potentially make Image Engine portable and capable of interacting with a wide range of external databases. As a first step towards this goal, the current prototype uses software agents to create real-time, dynamic links between the Image Engine database and the MARS clinical information system at UPMC.

MARS (**M**edical **AR**chival **S**ystem) is a text-based, clinical information system developed by John Vries and Russell Yount at UPMC. It contains approximately 4.5 million whole text, word-indexed, clinical records. These document records contain the full text of patient histories and physicals, operative and procedure notes from multiple clinical specialties, discharge summaries, laboratory results, and reports from the Pharmacy, Microbiology, Pathology and Radiology departments. The data stored on MARS includes 80% of all information generated at UPMC. More than 4,500 registered users retrieve an average of 20,000 reports each day.

Automatically retrieving image-related data from MARS requires four steps: (1) automatic creation of a database query from data stored in the Image Engine database; (2) automatic translation of this Image Engine query into a valid MARS database query; (3) execution of this MARS query and retrieval of clinical documents and (4) parsing and display of this MARS data within the Image Engine client. Steps one and four are handled largely by the Image Engine client application while steps two and three are the responsibility of the Image Engine MARS agent.

The MARS agent is an invisible, background application running on the same workstation as the Image Engine client. The agent shares a common, high-level, object-oriented view of clinical encounters with the Image Engine database. This shared data view is based on entities called Encounter Objects which are tokens representing a clinical encounter or set of encounters. For example, if the user selects a gastrointestinal endoscopy image in the Image Engine Thumbnail Browser and requests a MARS report for this image the client application automatically creates (from the Image Engine Object Database) a new Encounter Object representing the procedure that generated this image. This Encounter Object contains information about the patient the procedure was performed on, the procedure itself and the images generated by the procedure.

An Encounter Object is passed to the MARS agent which uses it to automatically create a MARS query, connect to the appropriate MARS database and retrieve clinical reports relevant to that clinical encounter. These reports are then inserted back into the original Encounter Objects and returned to the Image Engine Client for processing.

Because external database queries vary in complexity and therefore may take some time to complete (most agent-based MARS queries are completed within 10-15 seconds) Image Engine agents operate asynchronously and immediately return the client application to the user on receipt of an Encounter Object. When the agent is ready to return MARS data it inserts the modified Encounter Object into the client applications event queue for handling. The MARS agent uses simple rules to automatically expand MARS queries. For example, if the user requests a MARS report for an endoscopy image the agent will automatically expand this search to include any pathology or cytology reports related to the endoscopy procedure.

The MARS agent is a particular instance of a more general class of agents that we are developing as part of the Image Engine project. These agents share a common object-oriented view of clinical data with the Image Engine database but each agent class will be specialized to interact with a specific external database. For example, we are developing a Meta agent that communicates with NLM's Internet-based Metathesaurus server[13]. This agent may provide access to Metathesaurus data from within the Image Engine client.

The use of agents in this project reflects our belief that the migration towards component-based, data-centric computing offers many advantages when developing systems that share

and integrate clinical information. In this model, linking to additional external databases is possible by creating a new sub-class of agent specialized in translating Image Engine Encounter Objects into queries appropriate to the target database.

CLINICAL APPLICATIONS

We are currently testing Image Engine in three clinical environments at UPMC: Clinical Pathology, Gastroenterology and Medical Oncology. In Clinical Pathology we are focusing on issues related to digitizing, compressing, indexing, storing and retrieving both gross and microscopic pathology images. We have developed a pathology workstation using the Kodak DCS-420 high resolution color digital camera system for direct acquisition of digital microscopic pathology images. Clinicians in the Gastroenterology and Oncology test sites will identify pathology specimens (for example, biopsies obtained during gastrointestinal endoscopy or diagnostic oncology procedures). These pathology specimens will be digitized and added to the Image Engine database along with images obtained from related procedures. In Gastroenterology we are working with clinicians specializing in fiber optic endoscopy of the gastrointestinal and biliary tracts. These domains involve digital still and video images, pathology images and radiological imaging studies. In Medical Oncology we are exploring how one manages and integrates the wide range of images (including radiology, MRI, CT, pathology and clinical photography) that are used in the diagnosis, staging and treatment of patients with solid tumors.

ACKNOWLEDGMENT

The Image Engine project is supported by Biomedical Applications of High Performance Computing and Communications (BAA/HPCC) contract NO1-LM-4-3507 from the U.S. National Library of Medicine.

REFERENCES

1. MUN SK. Image management and communications (IMAC) in 1991. In: Heshiki, A Mun SK, (editors); Proceedings of the Second International Conference in Image Management and Communications in Patient Care.Computer Press; 1991;6-15.

2. GUR D, STRAUB WH, LIEBERMAN RH, GENNARI RC. Clinicians' access to diagnostic imaging information at an academic center: perceived impact on patient management. American Journal of Roentgenology 92 Apr;158(4):893-6

3. BEISECKER AE, BEISECKER TD. Patients Information-seeking behaviors when communicating with doctors. Med Care 1990; 28:19-28

4. ZINK S, JAFFE CC. Medical imaging databases. A National Institutes of Health workshop. Investigative Radiology 1993 Apr;28(4):366-72

5. GROSKY WI, MEHROTRA R. Image database management. IEEE Computer 1989;22:7-8.

6. CHOPLIN RH, BOEHME JM 2D, MAYNARD CD. Picture archiving and communication systems: an overview. Radiographics 1992 Jan;12(1):127-9.

7. BECKER SH, ARENSON RL. Costs and Benefits of Picture Archiving and Communication Systems. J Am Med Informatics Assoc. 1994;1:361-371

8. DAYHOFF RE, MALONEY DL, An integrated multidepartmental hospital imaging system: usage of data across specialties. Proceedings - the Annual Symposium on Computer Applications in Medical Care 1992;:30-4

9 LOWE, HJ, BUCHANAN BG, COOPER GF, VRIES JK. Building a Medical Multimedia Database System to Integrate Clinical Information: an Application of High Performance Computing and Communications Technology. Bull Med Libr Assoc 83(1) January 1995

10 YOUNT RJ, VRIES JK, COUNCIL CD: The Medical ARchival System: An information retrieval system based on distributed parallel processing. Information Processing and Management, 27:237-391,1991.

11 WALLACE GK. The JPEG Still Picture Compression Standard. Communications of the ACM, April 1991 (Vol. 34 no 4), 30-44

12 LE GALL D. MPEG: A Video Compression Standard for Multimedia Applications. Communications of ACM Vol 34;4 pp 46-63 (April 1991)

13 MCCRAY AT, RAZI A. The UMLS Knowledge Source Server. In Greenes R (Editor), Proceedings of the 8th. IMIA World Congress on Medical Informatics (Medinfo '95) Vancouver BC. July 1995